Finding periodic orbits of higher-dimensional flows by including tangential components of trajectory motion

Yang Wei Koh* and Kazuo Takatsuka†

Department of Basic Sciences, Graduate School of Arts and Sciences, The University of Tokyo, Komaba, 153-8902, Tokyo, Japan

(Received 6 June 2007; revised manuscript received 22 August 2007; published 17 December 2007)

Methods to search for periodic orbits are usually implemented with the Newton-Raphson type algorithms that extract the orbits as fixed points. When used to find periodic orbits in flows, however, many such approaches have focused on using mappings defined on the Poincaré surfaces of section, neglecting components perpendicular to the surface of section. We propose a Newton-Raphson based method for Hamiltonian flows that incorporates these perpendicular components by using the full monodromy matrix. We investigated and found that inclusion of these components is crucial to yield an efficient process for converging upon periodic orbits in high dimensional flows. Numerical examples with as many as nine degrees of freedom are provided to demonstrate the effectiveness of our method.

DOI: 10.1103/PhysRevE.76.066205 PACS number(s): 05.45.−a, 03.65.Sq, 45.50.Pk, 95.10.Eg

I. INTRODUCTION

Periodic orbits are important objects in the study of dynamical systems. In conservative systems, the winding numbers of periodic orbits in the integrable regime indicate the transition to chaos [1] as well as the topological structure of the phase space [2]. Periodic orbits can also provide information concerning the quantum mechanical behavior of classically chaotic systems: Gutzwiller's trace formula expresses the energy eigenvalues of chaotic systems in the semiclassical limit as a sum over all unstable periodic orbits [3]; also, Heller's discovery that eigenfunctions of the stadium billiard system are "scarred" along unstable periodic orbits indicates that periodic orbits might provide some clues concerning the eigenfunctions of chaotic systems [4,5]. The present authors have also discussed the amplitude-free mechanism of quantization of chaos based on the role of periodic orbits [6]. In dissipative systems, structural properties of strange attractors such as dimensions, Liapunov exponents, and topological entropy can be determined using unstable periodic orbits [7–9]. Even the existence of determinism in experimental time series can be revealed by the presence of periodicity in the experimental data [10].

Although both dissipative and conservative systems are described using ordinary differential equations, the different character of their periodic orbits gives rise to different techniques for their extraction. Techniques developed for each can in principle be carried over and applied in the other category, but there are usually some kinds of limitations involved. In dissipative systems, many fixed points are stable and a neighboring trajectory gets asymptotically closer to such a stable fixed point as time proceeds forward. However, in chaotic conservative systems this strategy cannot be applied because the periodic orbits are saddle fixed points [11]. Schmelcher and Diakonos proposed stabilizing all unstable manifolds of all periodic orbits [12–14]. Their method is based on a universal set of switching matrices that transforms the original system into one in which all the fixed points not only becomes stable but have global basins of convergence. The globally convergent property is important because the conventional Newton-Raphson method exhibits exponentially shrinking basins of attraction for unstable fixed points. However, for two dimensional systems, their method failed to stabilize parabolic motion, and this limitation carries over to higher dimensions [15]. In high dimensional conservative systems there are many unstable periodic orbits with mixed stability—the motion being parabolic along certain tangent manifolds and hyperbolic along others—and this restriction inherent in the stabilization method makes its usefulness limited. More importantly, the number of possible switching matrices one has to check scales as $N!2^N$ where N is the dimension of the system, making it a computationally expensive method to apply in higher dimensions. Crofts and Davidchack reduced this number drastically to 2^k where k is the dimension of the unstable subspace along the periodic orbit [15]. But the scaling is still exponential and in strongly chaotic regions of phase space with many unstable manifolds one still has to check through the many different possibilities. The stabilizing approach is hence difficult to apply in high dimensional conservative systems due to its inability to extract unstable periodic orbits with parabolic motions, high computational cost, and also the considerable mathematical analysis that might be involved.

An alternative approach more popular for conservative systems is the Newton-Raphson method [16]. One criticism of the Newton-Raphson method is that in chaotic systems the basin of attraction of the unstable saddle fixed points shrinks exponentially, rendering it difficult for the algorithm to succeed in converging upon the orbits [11,13,14]. More importantly, however, in the case of flows most studies have focused on first identifying a Poincaré surface of section upon which to construct the Jacobian used in the Newton-Raphson algorithm [11–14,17–19]. More precisely, one utilizes a Poincaré surface mapping of the form

$$\mathbf{x}' = (\mathbf{M}^{\mathrm{sur}})^p(\mathbf{x}), \tag{1}$$

where $\mathbf{M}^{\mathrm{sur}}$ denotes the mapping induced when the trajectory intersects once on the Poincaré surface, $\mathbf{x}$ and $\mathbf{x}'$ denote

*patrickkyw@mns2.c.u-tokyo.ac.jp
†Kaztak@mns2.c.u-tokyo.ac.jp

phase points constrained on the surface, and p the number of times the trajectory intersects the surface. To find a period p fixed point, one sets $\mathbf{x}'=\mathbf{x}$ in Eq. (1) and activates the Newton-Raphson algorithm to search for roots of the resulting equation. Although the Poincaré surface is a convenient tool for certain qualitative analyses, the dynamics of periodic orbits in flows of more than two dimensions are not necessarily well described by mappings defined on such surfaces of section. This is because a periodic orbit closes upon itself in the full phase space and not necessarily on the Poincaré surface; if the periodic orbit does not intersect near one's chosen surface of section, the Newton-Raphson procedure can fail to converge upon the periodic orbit. A standard way to treat this is to impose additional constraints in the mapping in Eq. (1) demanding that the solution lie on the Poincaré surface [35]. Such constraints, however, do not originate naturally from the phase flow of the system but are imposed externally in accordance with one's choice of Poincaré surface. Consequently, a solution is not guaranteed because a periodic orbit does not necessarily need to satisfy the one's choice of Poincaré surface.

To better address the shrinking basin as well as to circumvent the Poincaré surface problem, more sophisticated forms of Newton-Raphson methods called multipoint shooting methods have been proposed [19,20]. The idea is that by adopting more points along a periodic orbit, one can have an increased basin of attraction while at the same time achieving greater stability in the converging process. Indeed, it came to be realized that this approach to take more points along the periodic orbit, if taken to the extreme of abandoning the Poincaré surface altogether and directly seeking for the periodic orbit as an entire curve in the full phase space, yields a better approach to the problem. This gives rise to the so-called variational method approach to finding periodic orbits in flows [21].

In the variational method, one defines a certain cost function and a periodic orbit is obtained by seeking a curve that minimizes the cost function. However, this method requires one to solve a partial differential equation while at the same time manipulating an entire curve during the iteration process, making it possibly a computationally expensive process in high dimensions. By comparison, the original Newton-Raphson approach of using just the closest return does have certain merits such as ease of implementation and speed of calculation. Hence, in view of the advantages of doing away with the Poincaré surface, it would be interesting to see how the original Newton-Raphson method can be modified in a way that does not use a Poincaré surface while at the same time retaining its simplicity and utility.

In this study, we present a method to search for periodic orbits in Hamiltonian flows with many degrees of freedom (DOF). Our method utilizes the full monodromy matrix in the construction of the Jacobian that is used in the Newton-Raphson method, without resorting to the Poincaré surface of section approach. This allows matrix components corresponding to directions perpendicular to the Poincaré surface to participate in the Newton-Raphson iteration process. In other words, in contrast to Eq. (1), we consider a full phase space mapping of the form

$$\mathbf{z}' = \mathbf{M}^{\text{full}}(\mathbf{z}), \tag{2}$$

where $\mathbf{M}^{\text{full}}$ is the full phase space map, and $\mathbf{z}$ and $\mathbf{z}'$ denote points in phase space not constrained on any Poincaré surface [to be distinguished from $\mathbf{x}$ and $\mathbf{x}'$ in Eq. (1) which need to satisfy such constraints]. Our Newton-Raphson algorithm finds $\mathbf{z}$ such that $\mathbf{z}'=\mathbf{z}$ in Eq. (2). It is found that the additional information offered by these perpendicular components of the full monodromy matrix improves the convergence performance during the iteration process by allowing the trial orbit to deform more freely in the full phase space. We will refer to this aspect of our method as the full phase-space Newton-Raphson (FPSNR) algorithm to distinguish it from the standard NR algorithm based on the Poincaré section. In addition to the FPSNR algorithm, we realized that since we do not utilize a Poincaré surface, and also because the full monodromy matrix forces us to work with the full dimension of the phase space, we need a slightly different method of preparing initial conditions so that they can participate in the workings of the FPSNR algorithm. Hence we also present an algorithm that prepares close returns for this purpose.

In addition to being able to handle many-DOF systems, two other useful properties of our method are also worth mentioning. One is that it does not distinguish between types of periodic orbit, and hence can find periodic orbits regardless of their stability characteristics. In the numerical examples, eigenvalues of the monodromy matrices of the periodic orbits will be shown to justify our claim. The second property is that application of our method does not require any *a priori* knowledge of the structure or symmetries of the phase space of the system. Although such understandings are always preferable in the study of any system, the high dimensional and sometimes complicated forms of the potentials of realistic molecular systems usually makes it difficult to do this analytically. In such situations, our method is useful because the periodic orbits can be extracted directly, and can even be used in studying high dimensional phase space structure.

The paper is organized as follows. In Sec. II, we present the technical details of our method. A selection algorithm for preparing initial close returns is first described, followed by the full phase-space Newton-Raphson algorithm. In Sec. III we present examples of our method applied to systems with three, six, and nine DOF and discuss the results of our numerical calculations in detail. Section IV discusses and concludes the paper. Details concerning certain parameters which are used in the numerical calculations and sampling methods are discussed in the Appendixes.

II. METHOD

A. Selection algorithm: Preparing close returns in the full phase space

In the conventional approach of using a Poincaré surface of section, close returns (i.e., an initial guess of the periodic orbit) need to do two things: Intersect the Poincaré surface and be close to the starting point. When the phase space is very high dimensional and chaotic, these two conditions can

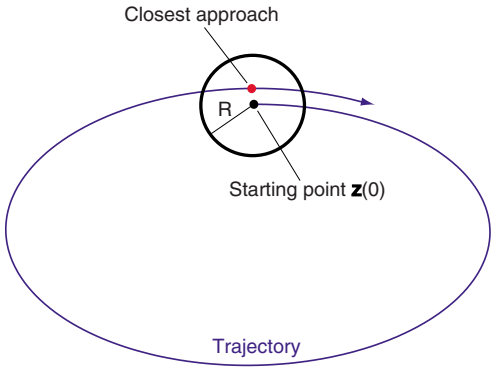

FIG. 1. (Color online) Small hyperspherical ball with radius R centered at the starting point of a trajectory. We want to determine t_{closest}, the time of closest approach to the starting point.

be difficult to satisfy simultaneously. First, a surface of section is a lower-dimensional plane embedded in the full space, and there are many ways of orienting the plane with respect to the full space. When the dimension of the full space is high and the surface of section a few dimensions smaller than the full space, a slightly different choice of orientation for the surface can easily lead to the trajectory not intersecting it. Hence a trajectory might have returned "close" to the starting point in some sense, but nevertheless disqualified because the Poincaré surface is not suitably oriented to intersect it. Near misses like these are costly when considered together with the second condition. When the phase space is strongly chaotic, a trajectory neighboring to an unstable periodic orbit may not survive being in its vicinity long enough to intersect the Poincaré surface as well as return close to the starting point. These discussions mean that in high dimensional chaotic phase space, the Poincaré surface of section approach makes it difficult to even prepare close returns. Besides, as discussed in the Introduction, even when close returns are successfully prepared, a Newton-Raphson procedure might not be able to converge onto a solution lying on the Poincaré surface.

Hence we adopt a slightly different way of preparing close returns. We would like to find as many periodic orbits as we can in an aimed subspace of the full phase space, such as a region of a given energy interval. Consider a small but nonzero volume of phase space surrounding a point lying on a periodic orbit; a trajectory initiating from a point in this neighboring region will lie close to the periodic orbit and return to the vicinity of its starting point just as the periodic orbit completes one transversal along its trajectory. We center a ball with the same dimension as the full phase space at the starting point and consider the portion of the trajectory reentering the ball as the close return (Fig. 1). In the following, we will describe in detail the procedure for picking out starting points with trajectories returning to the ball as well as determine the precise time at which the closest approach is made.

The first step is to pick a point in phase space and determine the usable lifetime, denoted t_{cut}, of the stability matrix along the trajectory issuing from this point [22–24]:

$$\mathbf{M}'(t) = \begin{pmatrix} \frac{\partial \mathbf{p}(t)}{\partial \mathbf{p}(0)} & \frac{\partial \mathbf{p}(t)}{\partial \mathbf{q}(0)} \\ \frac{\partial \mathbf{q}(t)}{\partial \mathbf{p}(0)} & \frac{\partial \mathbf{q}(t)}{\partial \mathbf{q}(0)} \end{pmatrix}. \tag{3}$$

When the system is chaotic, machine roundoff errors accumulate in $\mathbf{M}'(t)$ during the course of integrating the trajectory. Since we use the monodromy matrix in the full phase-space Newton-Raphson algorithm we need to ensure the numerical accuracy of $\mathbf{M}'(t)$. One criterion would be to check the determinant of $\mathbf{M}'(t)$ and stop integration when it deviates too much from unity. Appendixes A and B discuss this and the picking of phase points in more detail.

After determining t_{cut}, we pinpoint the time, t_{closest}, when the trajectory makes its closest approach to the starting point. Center a small hyperspherical ball with radius R at the starting point of a trajectory (see Fig. 1). For the portion of the trajectory $t < t_{\text{cut}}$, compute the distance $D(t)$ between the starting point $\mathbf{z}(0)=[\mathbf{p}(0),\mathbf{q}(0)]$ and the points along the trajectory $\mathbf{z}(t)=[\mathbf{p}(t),\mathbf{q}(t)]$ using the Euclidean norm [36]

$$D(t) = \sqrt{[\mathbf{z}(t) - \mathbf{z}(0)]^2}. \tag{4}$$

If $D(t) > R$ for all $t < t_{\text{cut}}$, the starting point is discarded since it gives no close return. If $D(t) < R$ for some $t < t_{\text{cut}}$, we determine, for the portion of the trajectory lying inside the ball, the point where $D(t)$ is a minimum within that segment and set that t as t_{closest}. In other words, we determine the time at which the trajectory makes a closest approach to its starting point [37]. The starting point $[\mathbf{p}(0),\mathbf{q}(0)]$ will be our first guess of a periodic orbit and t_{closest} is set to be the period T of that periodic orbit. The FPSNR algorithm is then applied to them to get a more accurate periodic orbit. In the course of applying the FPSNR algorithm, t_{closest} is held fixed while $[\mathbf{p}(0),\mathbf{q}(0)]$ is gradually adjusted until a periodic orbit is obtained. It is possible that one starting point can give rise to many different t_{closest}. In the case of unstable trajectories, different t_{closest} for the same starting point may lead the FPSNR algorithm to converge onto different periodic orbits, so one should apply the FPSNR algorithm to all of the different t_{closest} in these cases.

To summarize: Pick a point in phase space and determine the t_{cut} for that point. If the trajectory returns near its starting point before t_{cut}, determine t_{closest} for the segment of that trajectory within the ball. The starting point and corresponding t_{closest} obtained this way will serve as the initial guesses for the FPSNR algorithm.

Figure 2 summarizes the selection algorithm in the form of a flow chart.

B. Full phase-space Newton-Raphson algorithm using the full monodromy matrix

We now present the full phase-space Newton-Raphson algorithm that locates a periodic orbit using the set of points and times obtained from the selection algorithm. Let us denote t_{closest} as T because it is the period of the periodic orbit we are going after. Referring to Fig. 3, we see that if we integrate the trajectory starting from $\mathbf{z}$ for time T, we arrive

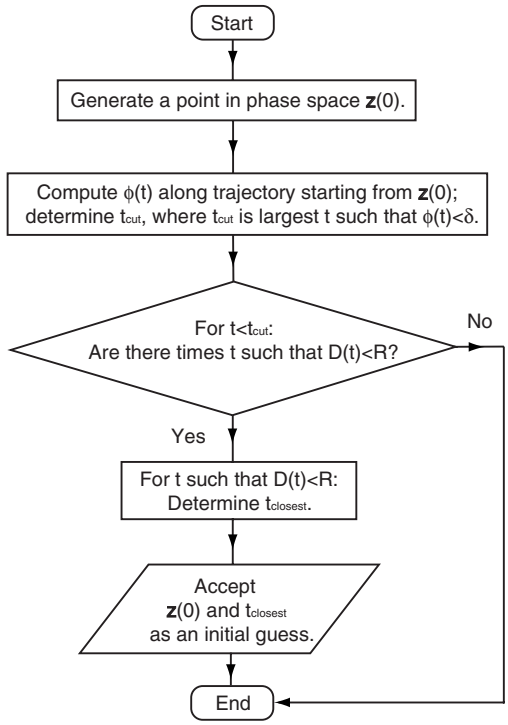

FIG. 2. Flow chart for the selection algorithm. $\phi(t)$ and δ are described in Appendix B.

under the full phase space map $\mathbf{M}$ at $\mathbf{M}(\mathbf{z},T)$ such that there is a vector $d\mathbf{z}$ pointing from $\mathbf{z}$ to $\mathbf{M}(\mathbf{z},T)$, that is $d\mathbf{z}=\mathbf{M}(\mathbf{z},T)-\mathbf{z}$. Our proposed FPSNR algorithm computes a correction vector $\mathbf{c}$ and add it to $\mathbf{z}$ to obtain a new initial condition $\mathbf{z}_{\text{new}}$ such that the length of $d\mathbf{z}_{\text{new}}$ is smaller than that of $d\mathbf{z}$ (see Fig. 3 for the definitions of these vectors); $\mathbf{z}$ is then repeatedly adjusted until the magnitude of $d\mathbf{z}_{\text{new}}$ is below some small tolerance length.

In seeking for the fixed points of a mapping $\mathbf{F}(\mathbf{x})$, one is looking for roots of the equation

$$\mathbf{G}(\mathbf{x}) \equiv \mathbf{F}(\mathbf{x}) - \mathbf{x} = 0. \tag{5}$$

Suppose we have an initial guess $\mathbf{x}$ of the solution for Eq. (5) and wish to compute a next $\mathbf{x}_{\text{new}}$ which is a better solution, the Newton-Raphson iteration prescribes the following iteration scheme

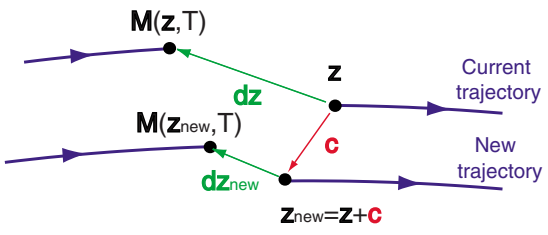

FIG. 3. (Color online) Closeup schematic diagram of the neighboring region of a starting point $\mathbf{z}$. Under the full phase space map $\mathbf{M}$, the trajectory moves around phase space and arrives at $\mathbf{M}(\mathbf{z},T)$ after time T. Adding the vector $\mathbf{c}$ to $\mathbf{z}$ will result in a $d\mathbf{z}_{\text{new}}$ with a smaller magnitude.

$$\mathbf{x}_{\text{new}} = \mathbf{x} + [\mathbf{G}'(\mathbf{x})]^{-1}\mathbf{G}(\mathbf{x}), \tag{6}$$

where $\mathbf{G}'(\mathbf{x})$ is the Jacobian of $\mathbf{G}$ evaluated at $\mathbf{x}$. When finding fixed points on a Poincaré surface, one substitutes the Poincaré map $\mathbf{F}(\mathbf{x})$ into Eq. (6) via Eq. (5) to get

$$\mathbf{x}_{\text{new}} - \mathbf{x} = [\mathbf{F}'(\mathbf{x}) - \mathbf{1}]^{-1}[\mathbf{F}(\mathbf{x}) - \mathbf{x}], \tag{7}$$

where $\mathbf{x}$ and $\mathbf{x}_{\text{new}}$ are constrained on the Poincaré surface. The correction $\mathbf{x}_{\text{new}}-\mathbf{x}$ thus obtained is then added to $\mathbf{x}$ to get $\mathbf{x}_{\text{new}}$.

As argued in the Introduction, Eq. (7) using the Poincaré map with points constrained on the surface is too restrictive and may result in no solution. Instead, we propose using the full monodromy matrix in Eq. (6) to arrive at the following iteration scheme:

$$\mathbf{z}_{\text{new}} - \mathbf{z} = [\mathbf{M}'(\mathbf{z},T) - \mathbf{1}]^{-1}[\mathbf{M}(\mathbf{z},T) - \mathbf{z}], \tag{8}$$

where $\mathbf{M}$ is the full phase space map and $\mathbf{M}'(\mathbf{z},T)$ is the full monodromy matrix calculated along a trajectory starting from $\mathbf{z}$ for time T (both obtained from the selection algorithm). The key is that in using the full monodromy matrix $\mathbf{M}'$, $\mathbf{z}_{\text{new}}$ is no longer constrained on any Poincaré surface. Whereas a mapping defined on a Poincaré surface can only indicate the direction to move within the confines of the surface, the direction provided by $\mathbf{M}'$ can allow—with additional perpendicular components leading out of the surface—for better adjustments in accordance with the phase flow of the system and hence approach the periodic orbit more naturally.

Denoting $\mathbf{z}_{\text{new}}-\mathbf{z}$ as $\Delta\mathbf{z}$ and $\mathbf{M}(\mathbf{z},T)-\mathbf{z}$ as $d\mathbf{z}$, Eq. (8) can be rewritten as

$$\Delta\mathbf{z} = [\mathbf{M}'(\mathbf{z},T) - \mathbf{1}]^{-1}d\mathbf{z}. \tag{9}$$

Equation (9) is a linear equation and seems to be readily solvable for $\Delta\mathbf{z}$ using the standard techniques [25]. However, as is well known, at least a pair of eigenvalues of the matrix $\mathbf{M}'$ approach unity as the trajectory becomes close to a periodic orbit, and thereby the matrix $\mathbf{M}'-\mathbf{1}$ will turn out to be singular. A standard treatment to avoid such a singularity would be to remove the matrix components corresponding to tangential motion along the orbit. Indeed, methods based on the Poincaré surface of section actually use this procedure, that is, accounting for the periodic motion in the transversal directions only. However, as discussed in the Introduction, this can lead to failure to arrive at the periodic orbit aimed at. It is crucial to include the tangential components to deform and guide a trial trajectory along the right direction to make it close itself. To compromise these conflicting demands, we here define the *correction vector* $\mathbf{c}$ as

$$\mathbf{c} \equiv -\epsilon|d\mathbf{z}|\frac{\Delta\mathbf{z}}{|\Delta\mathbf{z}|}, \tag{10}$$

where ϵ is a small fraction and $|\cdot|$ means the magnitude of the vector. What Eq. (10) means is that we rescale the $\Delta\mathbf{z}$ we obtained from solving Eq. (9) to only a fraction ϵ of the length of the vector $d\mathbf{z}$ and reverse its sign. Inclusion of $|d\mathbf{z}|$ as a scaling factor is critical in this rescaling procedure. The calculation of $\mathbf{c}$ via Eqs. (9) and (10) constitutes a simple damped Netwon-Raphson iteration procedure for locating

the periodic orbit. The parameter ϵ can further be optimized to achieve faster convergence. However, we would not undertake this task here but simply chose ϵ heuristically (see Appendix C).

After $\mathbf{c}$ is obtained, it is added to $\mathbf{z}$ to give a new initial condition:

$$\mathbf{z}_{\text{new}} = \mathbf{z} + \mathbf{c}. \tag{11}$$

The vector $d\mathbf{z}_{\text{new}}$ (see Fig. 3) is computed using $\mathbf{z}_{\text{new}}$. If its magnitude is smaller than some tolerance limit γ, $\mathbf{z}_{\text{new}}$ is accepted as an initial condition of a periodic orbit; otherwise, the above process is repeated until $|d\mathbf{z}_{\text{new}}| < \gamma$. We remark here that $\mathbf{c}$ does not conserve energy and $\mathbf{z}_{\text{new}}$ in general has a different energy from $\mathbf{z}$.

It is mathematically true that the full use of the monodromy matrix in Eq. (9) leads to a divergence at the exact point where the closure of an orbit is rigorously completed. To avoid this, one may use other numerical techniques such as the singular value decomposition method to complete the closure after a sufficiently good convergence has been attained with the use of Eq. (9). Nevertheless, we did not undertake this additional process in this paper, since the procedure of Eq. (9) already gives a monotonic and very good convergence to an exact periodic orbit as shown in the next section. Moreover, once an approximate periodic orbit is found numerically, one can just make repeated use of the first circuit for other calculations that require multiple transversals along the same periodic orbit, without integrating the trajectory beyond the closure point. Therefore, the simple NR procedure we have proposed above is good enough in practice, and no further pursuit for the exact periodic orbit has been attempted.

Figure 4 summarizes the FPSNR algorithm in the form of a flow chart.

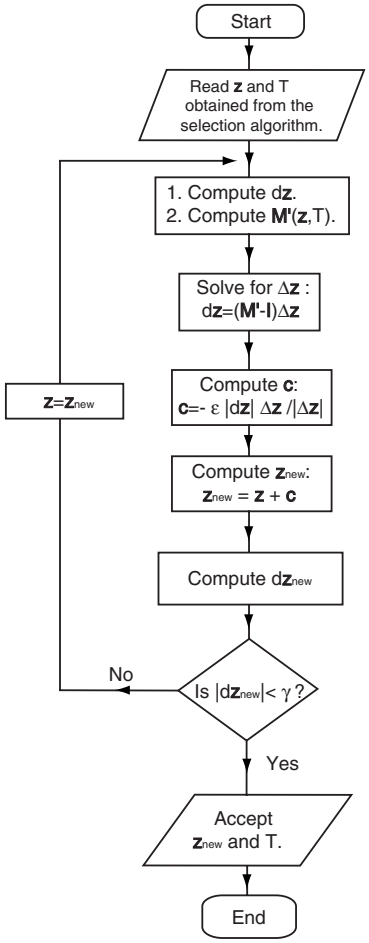

FIG. 4. Flow chart for the full phase-space Newton-Raphson algorithm. T here is equivalent to t_{closest} from Fig. 2.

III. NUMERICAL EXAMPLES

This section presents the numerical results of applying our method to search for periodic orbits in some Hamiltonian systems. Our presentation is meant to highlight some aspects of our method. To demonstrate its applicability to high dimensional systems, we tested our method on the generalized Hénon-Heiles systems with six and nine DOF [26,27]. We also test it on two other systems with three DOF. These systems and their DOF are summarized in Table I.

We would also like to emphasize that our method is able to find periodic orbits of any stability characteristics. To illustrate this, we compute the eigenvalues λ of the monodromy matrix corresponding to each periodic orbit and plot the eigenvalues schematically on the complex plane to give an idea of their distribution.

Also, the selection algorithm and the FPSNR algorithm are very effective in preparing good close returns and successfully locating the periodic orbit, respectively. Let us consider some numbers. The first one is N_{grid}. This is the number of grid points used in the selection algorithm (Appendix A). The second number is N_{guess}. This is the number of close returns obtained from the grid points via the selection algorithm. The third number is N_{po}. This is the total number of periodic orbit conditions obtained from the initial guesses after applying the FPSNR algorithm. The results for N_{grid} for each system are summarized in Table II while those for N_{guess} and N_{po} are summarized in Table III.

The effectiveness of the two algorithms can be measured by considering two ratios r_1 and r_2. The first one is defined as

$$r_1 = \frac{N_{\text{guess}}}{N_{\text{grid}}}. \tag{12}$$

r_1 is a measure of the effectiveness of the selection algorithm in preparing close returns for use in the FPSNR algorithm. A

TABLE I. Types of systems used in our numerical calculations, their names, and their DOF.

Type of system	Name	DOF
Generalized	A	3
Hénon-Heiles	B	6
	C	9
$x^2y^2z^2$ potential	D	3

TABLE II. Summary of how the grid is taken for the systems investigated. The region of phase space in which the grid is laid out is given in the second column. Spacing refers to the distance between grid points. N_{grid} is the total number of grid points.

System	Grid	Spacing	N_{grid}
A	$3 \geq p_x, p_y, p_z \geq 0$, $3 \geq x, y, z \geq 0$	Along all axes: 0.3	$11^6 = 1\,771\,561$
B	$x_1 = \cdots = x_6 = 0$, $\pi > \psi_1, \ldots, \psi_4 > 0$, $2\pi \geq \psi_5 > 0$	Angular spacing: $\frac{\pi}{10}$	$9^4 \times 20 = 131\,220$
C	$x_1 = \cdots = x_9 = 0$, $\pi > \psi_1, \ldots, \psi_7 > 0$, $2\pi \geq \psi_8 > 0$	Angular spacing: $\frac{\pi}{6}$	$5^7 \times 12 = 937\,500$
D	$2 \geq p_x, p_y, p_z \geq -2$,	Along p_x, p_y, p_z axes: 0.4	$11^3 \times 11^3 = 1\,771\,561$
	$2 \geq x, y, z \geq 0$	Along x, y, z axes: 0.2	

small r_1 indicates that the selection process had been helpful. The ratio r_2 is defined as

$$r_2 = \frac{N_{\text{po}}}{N_{\text{guess}}}. \tag{13}$$

It is the success probability of the FPSNR algorithm in converging onto a periodic orbit. A r_2 close to unity indicates that the FPSNR algorithm has a high success rate. The relatively high values for r_2 we obtained in our numerical calculations show that the FPSNR algorithm is indeed effective in locating periodic orbits. The values for r_1 and r_2 for each of the systems we investigated are also given in Table III.

Lastly, we think that it is of interest to show the extent to which $|d\mathbf{z}|$ can be reduced using the FPSNR algorithm. As a measure of the size of $|d\mathbf{z}|$ of a periodic orbit, we compare it with the perimeter of the periodic orbit itself. The perimeter of the periodic orbit P is computed by summing up the Euclidean distances between consecutive points along the periodic orbit; that is,

TABLE III. Summary of algorithm parameters and some numerical results for the systems investigated. R is the radius of the ball employed in the selection algorithm (Fig. 1). N_{guess} is the number of initial guesses obtained from N_{grid} (Table II). ϵ is the scaling factor for the correction vector given in Eq. (10). γ is the comparison tolerance length used in the FPSNR algorithm (Fig. 4). N_{po} is the number of periodic orbits obtained by the FPSNR algorithm from N_{guess}. r_1 and r_2 are defined by Eqs. (12) and (13), respectively. The N_{po} for system C is in parentheses because we applied the FPSNR algorithm to only four of the N_{guess} initial guesses.

System	R	N_{guess}	ϵ	γ	N_{po}	r_1	r_2
A	0.4	108 019	0.05	0.0025	20 459	0.061 0	0.189
B	0.5	55	0.025	0.001	44	0.000 419	0.8
C	0.5	938	0.1	0.0001	(4)	0.001 00	
D	0.4	728	0.05	0.01	619	0.000 411	0.850

$$P = \sum_{i=1}^{n} \sqrt{[\mathbf{z}(i+1) - \mathbf{z}(i)]^2}, \tag{14}$$

where $\mathbf{z}(i)$ is the point on the periodic orbit at the ith time step, and the periodic orbit makes its closest approach to its starting point at the nth time step. We measure the size of $|d\mathbf{z}|$ of each periodic orbit using Δ, defined as

$$\Delta \equiv \frac{|d\mathbf{z}|}{P}. \tag{15}$$

The Δ for the periodic orbits presented are calculated and presented in the tables.

In the following, we will present and discuss the numerical results for the systems we investigated.

A. Generalized Hénon-Heiles systems

We first applied our method to generalized Hénon-Heiles systems [26,27]. The Hamiltonian is given by

$$H_N(p_1, \ldots, p_N, x_1, \ldots, x_N) = \frac{1}{2}\sum_{i=1}^{N} p_i^2 + \frac{1}{2}\sum_{i=1}^{N} x_i^2 + \alpha \sum_{i=1}^{N-1} \left(x_i^2 x_{i+1} - \frac{x_{i+1}^3}{3} \right), \tag{16}$$

where $\alpha = 0.111\,803$ is a nonlinearity parameter and N is the DOF of the system. Low-dimensional Hénon-Heiles systems have been popular subjects in the study of quantum and semiclassical mechanics [28,29], partly because they offer an interesting mixture of regular and chaotic motions. Recently, higher-dimensional versions have been used as models in several multidimensional semiclassical quantization studies [26,27]. We have chosen to investigate this popular system using the same value for the nonlinearity parameter.

We applied our method to the cases of $N=3$, 6, and 9. In the following, we first discuss in detail the case for $N=3$.

This system is high dimensional enough to be nontrivial but also low dimensional enough for us to give a comprehensive presentation of the myriad of orbit topologies and the stability characteristics of the periodic orbits found using our method. We then go on to the cases of $N=6$ and 9 DOF to demonstrate the effectiveness of our method in higher-dimensional situations.

1. N=3

Our first example is for the case of $N=3$. Let us relabel the configuration coordinates x_1, x_2, and x_3 as x, y, and z, respectively. We first describe the calculation of N_{grid}, N_{guess}, and N_{po}. The grid is laid out in the region of phase space given in Table II. Along the direction of each axis the distance between consecutive grid points is 0.3, giving $N_{\text{grid}}=11^6=1\,771\,561$. Letting $R=0.4$ and considering only the first closest approach, the selection algorithm returns $N_{\text{guess}}=108\,019$ initial guesses. Using $\epsilon=0.05$ and $\gamma=0.0025$ in the FPSNR algorithm, one finally obtains $N_{\text{po}}=20\,459$ periodic orbits. The values of these results are also summarized in Tables II and III.

In the following, ten periodic orbits labeled A1 to A10 are selected for discussion. Figure 5 shows for A1 to A6 the trajectories of the orbits in configuration space together with schematic plots of the eigenvalues of their monodromy matrix on the complex plane. Those of A7 to A10 are shown in Fig. 6. Table IV also gives the energy, perimeter, and Δ of each periodic orbit.

We first discuss the stability characteristics of the periodic orbits. A1 to A6 are stable while A7 to A10 are unstable. A1, A4, and A6 have three complex conjugate pairs of eigenvalues on the unit circle; A2 and A5 have two such pairs on the unit circle and the last pair splits from unity along the real axis; A3 has only one complex conjugate pair on the unit circle and the remaining four form a four-tuple near unity. Hence, even if there exist unit eigenvalues in the monodromy matrix, our simple damped NR iteration scheme can still successfully converge onto the periodic orbit.

We next discuss the unstable periodic orbits. A7 is a periodic orbit with a mixture of both stable and unstable motions: it has one pair of real, positive eigenvalues indicating hyperbolic motion, and a second pair lying on the unit circle indicating elliptic motion. A8 has two pairs of eigenvalues lying on the real axis and the last pair splitting from unity along the unit circle. Lastly, A9 and A10 are so-called loxodromic periodic orbits, with four of their eigenvalues forming a four-tuple. Hence, the FPSNR algorithm is also capable of locating unstable periodic orbits of any stability characteristic.

Lastly, we present the values of Δ (and indirectly, $|d\mathbf{z}|$) for A1 to A10 to show that they are small and hence justify that they can reasonably be considered as accurate periodic orbits. The Δ for each of the periodic orbits is listed in Table IV, and we can see that $|d\mathbf{z}|$ is small compared to the perimeter of the entire trajectory. Alternatively, one can obtain a rough value of $|d\mathbf{z}|$ by multiplying Δ to P and then compare its magnitude to the scales given on the axes of the trajectory plots to get a rough idea of their relative size. Either way, we see that the FPSNR algorithm is indeed capable of reducing $|d\mathbf{z}|$ to the extent that A1 to A10 can be considered as periodic orbits.

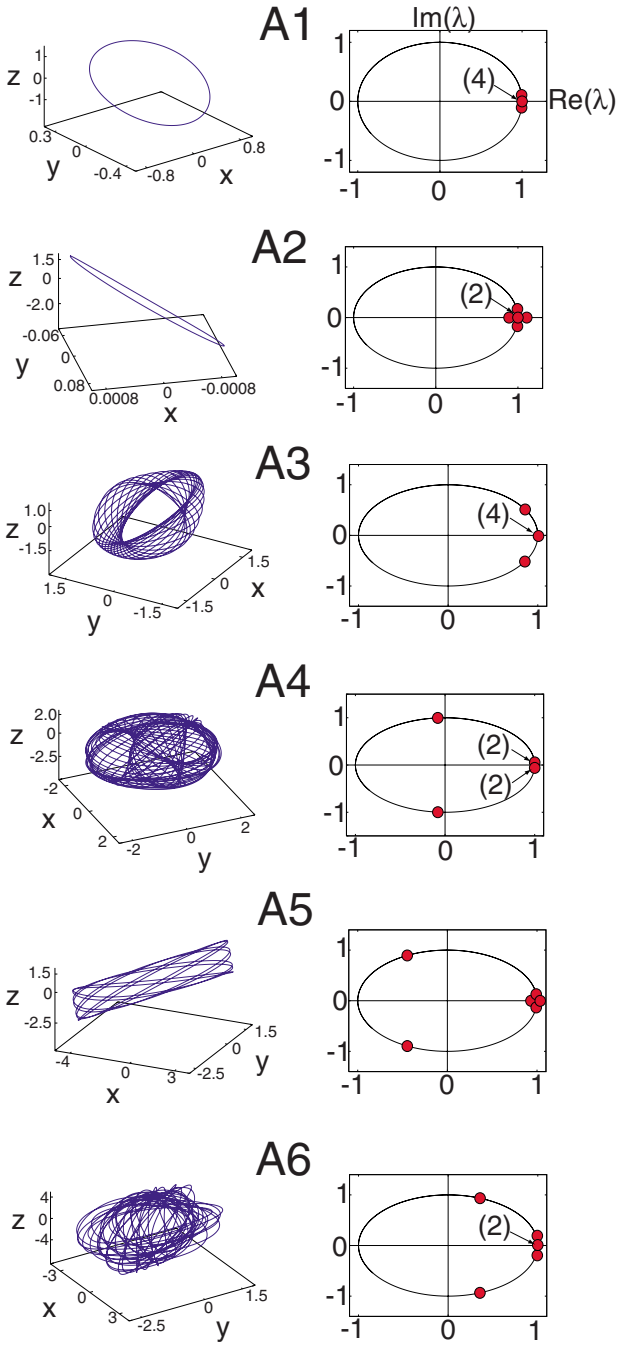

FIG. 5. (Color online) Stable periodic orbits A1 to A6 of the generalized three-DOF Hénon-Heiles system. The left-hand panels show the trajectory plots of the periodic orbits in configuration space. The corresponding panels on the right side show schematic plots of the monodromy matrix eigenvalues on the complex plane. As some eigenvalues may be very close to each other and hence difficult to distinguish in the plots, we further indicate these cases by numbers in parentheses.

2. High-dimensional cases N=6,9

As the number of grid points scales exponentially with dimension, when dealing with high dimensional systems it is

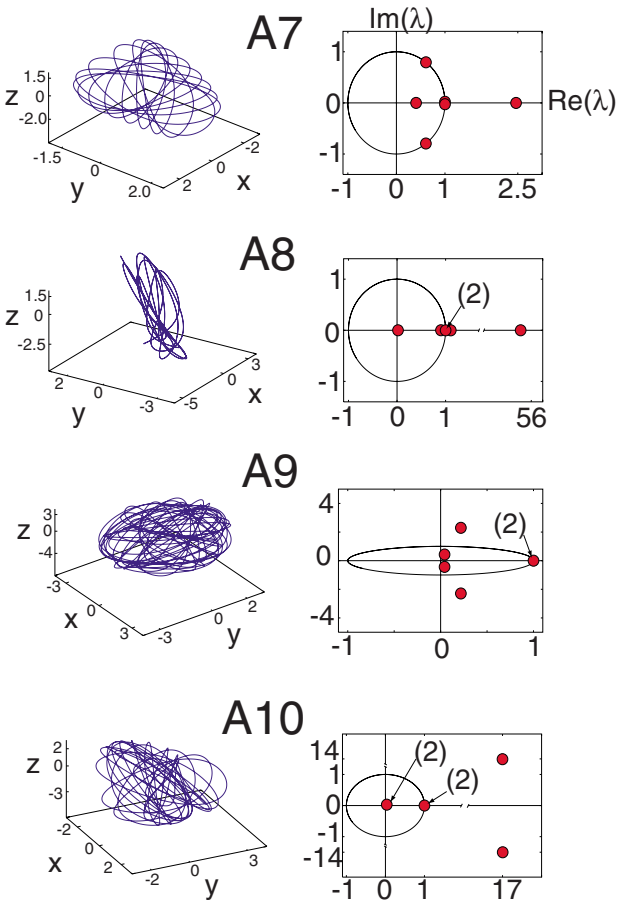

FIG. 6. (Color online) Unstable periodic orbits A7 to A10 of the generalized three-DOF Hénon-Heiles system.

sometimes necessary to focus the selection algorithm on just those grid points falling within the energy range one is interested in studying. For the six and nine degrees of freedom system, the grid is laid out using hyperspherical coordinates and only grid points within a certain energy shell is subjected to the selection algorithm. Details concerning how the grid is taken using these coordinates are described in Appendix A.

TABLE IV. Numerical data for periodic orbits A1 to A10 of the three-DOF generalized Hénon-Heiles system. P is the perimeter of the periodic orbit. Δ is defined in Eq. (15).

System A			
Name of periodic orbit	Energy	P	Δ
A1	0.89	8.40	1×10^{-5}
A2	1.40	10.55	9×10^{-6}
A3	2.30	243.12	4×10^{-7}
A4	4.86	701.37	1×10^{-7}
A5	7.86	183.27	6×10^{-6}
A6	8.99	854.20	1×10^{-7}
A7	4.44	206.66	1×10^{-5}
A8	9.55	446.00	6×10^{-7}
A9	8.99	720.33	1×10^{-7}
A10	6.04	416.84	2×10^{-7}

Table II summarizes the number of grid points obtained. Table III summarizes the parameters used in the selection and FPSNR algorithms as well as the N_{grid} and N_{po} obtained.

We first discuss an important role played by the selection algorithm in dealing with higher-dimensional systems. We chose to lay out the grid in relatively high-energy regions of phase space, and many of the trajectories dissociate to infinity. Hence, on one hand, the number of phase points is very large for many DOF systems because of the dimensionality of the phase space involved; on the other hand, the actual number of close returns is only a very small fraction of all the phase points. The small values of 0.000 419 and 0.001 of r_1 that we obtained for $N=6$ and 9 indicate that the selection algorithm performs an important role of preselection in helping to locate periodic orbits in high-dimensional systems. Without weeding out the bad candidates, it would be very difficult to rely solely on the FPSNR algorithm alone to locate the periodic orbits in these systems.

In the following, we will present some of the periodic orbits we found for these two systems. Five periodic orbits B1 to B5 from system B are presented in Figs. 7 and 8. Four orbits C1 to C4 from system C are presented in Figs. 9 and 10. The energies, perimeters, and Δ of these periodic orbits are also given in Table V.

B1 to B3 are quasistable periodic orbits. We see from B3 that even for a six-DOF system our method can locate long and complex periodic orbits. Unstable periodic orbits can also be found, such as B4 and B5. Roughly the same situation occurs in system C, where in the high-energy regime most of the periodic orbits we obtained are unstable, short ones such as C1 to C3; nevertheless, it is still possible to get relatively long and complex—yet still unstable—ones such as C4.

Even though our NR algorithm is not optimized, it is capable of very quickly and effectively reducing the $|d\mathbf{z}|$ for high-dimensional periodic orbits, as indicated by the small Δ listed in Table V for all the periodic orbits. To give a further idea of the convergence process, Fig. 11 shows the graph of how the $|d\mathbf{z}|$ of the sequence of trajectories leading to the periodic orbit C2 decreases with each application of the FPSNR algorithm. We began with an initial guess trajectory whose $|d\mathbf{z}|$ is approximately 0.45 and this $|d\mathbf{z}|$ was reduced to less than 0.0001 with less than 300 repetitions of the algorithm.

B. Strongly chaotic system: $x^2y^2z^2$ potential

We have demonstrated that our method is applicable to the Hénon-Heiles type systems with many DOF. In this subsection, we will apply it to a strongly unstable system to show that our method can find periodic orbits even in strong chaos. We chose the following three-DOF system

$$H(p_x,p_y,p_z,x,y,z)=\frac{1}{2}(p_x^2+p_y^2+p_z^2)+\frac{1}{2}x^2y^2z^2. \quad (17)$$

This system is strongly chaotic. In finding periodic orbits in strong chaos, a key role is played once again by the selection algorithm in preparing good initial guesses. In such systems, trajectories that get a chance to return to near their starting

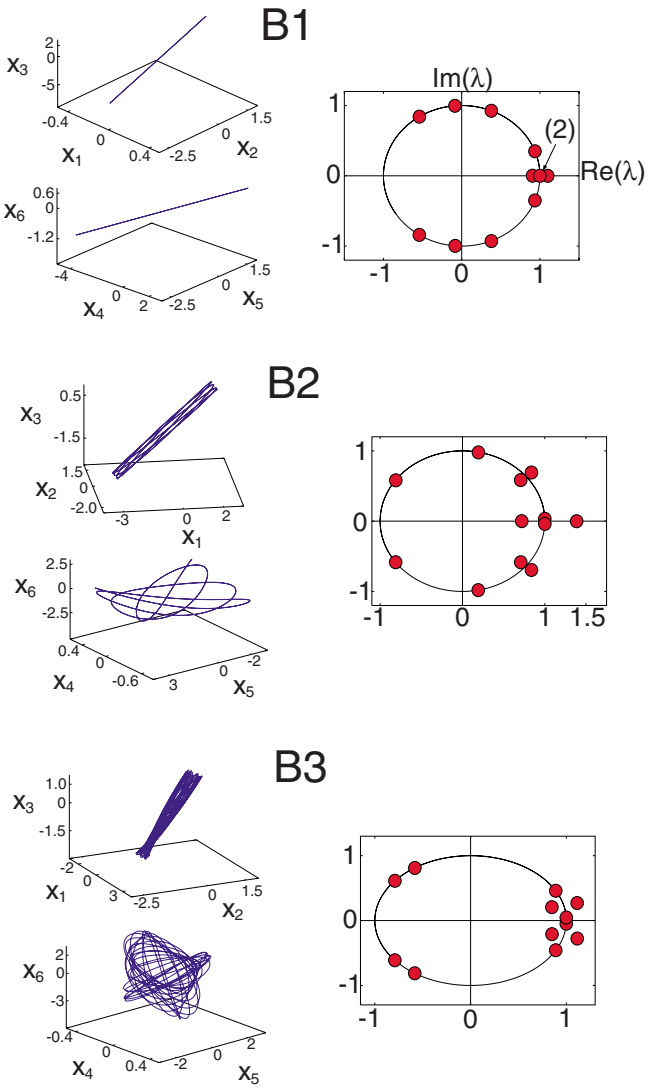

FIG. 7. (Color online) Quasistable periodic orbits B1 to B3 of the six-DOF generalized Hénon-Heiles system. Projections of each orbit onto the x_1-x_2-x_3 and the x_4-x_5-x_6 spaces are shown.

points are very rare, and it is important to first locate these with the selection algorithm before applying the FPSNR algorithm. This can be discerned by the smallness of the ratio $N_{\mathrm{po}}/N_{\mathrm{grid}} \approx 0.0003$ we obtained (Table III). If one were to blindly select points from phase space and subject them straight away to the FPSNR algorithm, the chance of locating any periodic orbit would be extremely small.

We have selected three periodic orbits, labeled D1 to D3, to present in Fig. 12. Since this system exhibits scaling invariance, we leave out the scales on the axes in the trajectory plots. (See [32] for quantization of chaos making use of the scale invariance of the Hamiltonian.) The energy, perimeter, and Δ for each periodic orbit are given in Table VI.

A large proportion of the periodic orbits we obtained are confined to planes in configuration space. This is because the strongly unstable nature of the system makes it difficult for any trajectory to come back to its starting point once it begins to perform three-dimensional motion. Nevertheless, we still succeeded in locating many periodic orbits which are

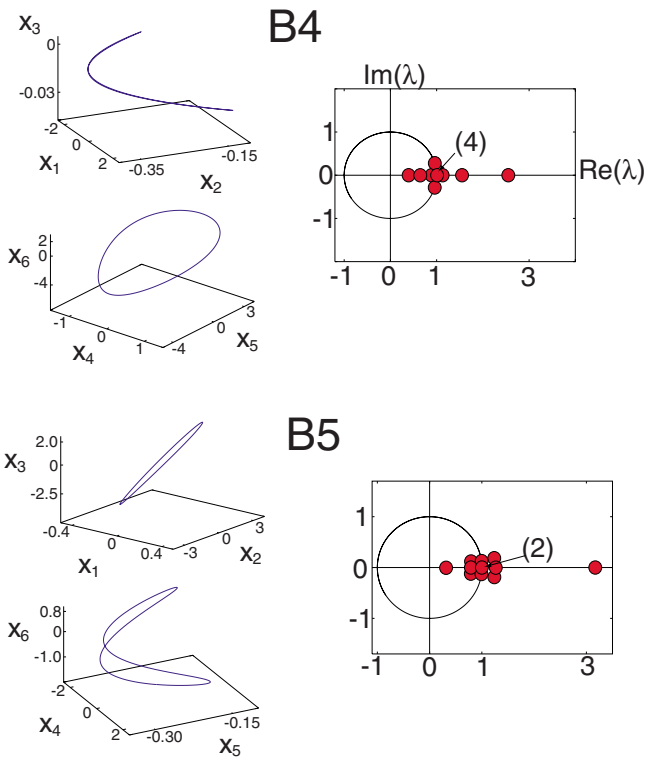

FIG. 8. (Color online) Unstable periodic orbits B4 and B5 of the six-DOF generalized Hénon-Heiles system.

three-dimensional curves. D1 to D3 are examples of such orbits. Also, the monodromy matrix eigenvalues of these periodic orbits indicate that they are very unstable, and the fact that we can find them demonstrates that our method is effective in handling even strongly chaotic systems.

IV. DISCUSSION AND CONCLUSION

As mentioned at the beginning of Sec. III, the ratios r_1 and r_2 provide measures for the effectiveness of the selection and FPSNR algorithms, respectively. Strictly speaking, r_1 and r_2 may not be the most stringent way to assess the performance of our method, but we think that it gives a rough idea of how well the selection and FPSNR algorithm works in practice. The small r_1 we obtained for all the systems investigated—especially in strongly chaotic situations such as those of B, C, and D—shows that the selection algorithm plays an important role in directing and focussing the FPSNR algorithm on only those important points in phase space.

The low number of r_2 we obtained for system A demands some explanation. In this system, many of the initial guesses give rise to trajectories lying in the region of phase space where there are channels of dissociation. In the course of adjusting such a trajectory using the convergence algorithm, the trajectory might fall into one of these dissociation channels and escape to infinity. As a caveat when using the FPSNR algorithm, one should always monitor the monodromy matrix of the trajectory during the NR iteration process so as to be aware whenever such dissociation cases occur.

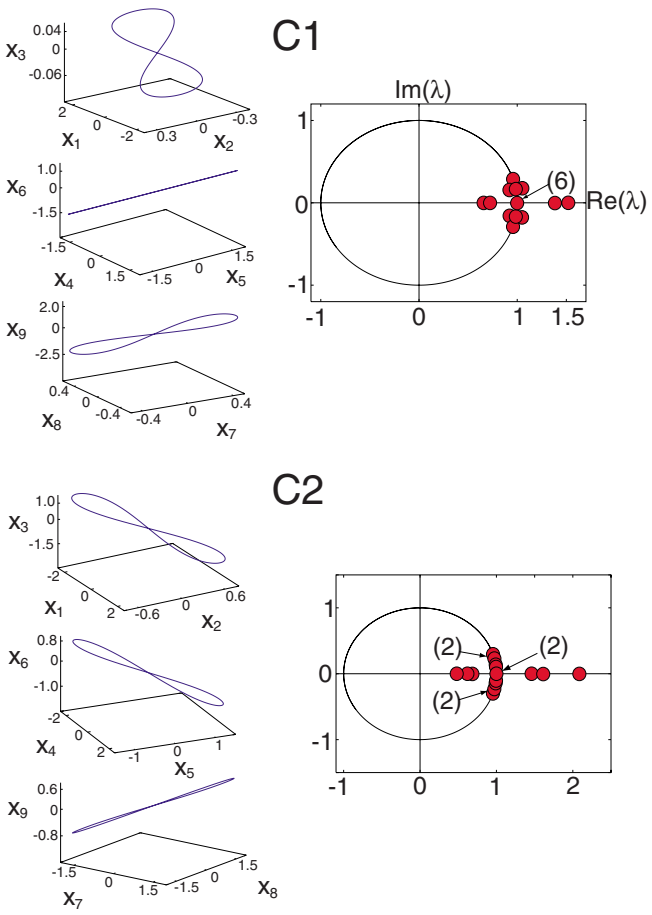

FIG. 9. (Color online) Periodic orbits C1 and C2 of the nine-DOF generalized Hénon-Heiles system. Projections of each orbit onto the x_1-x_2-x_3, the x_4-x_5-x_6, and the x_7-x_8-x_9 spaces are shown.

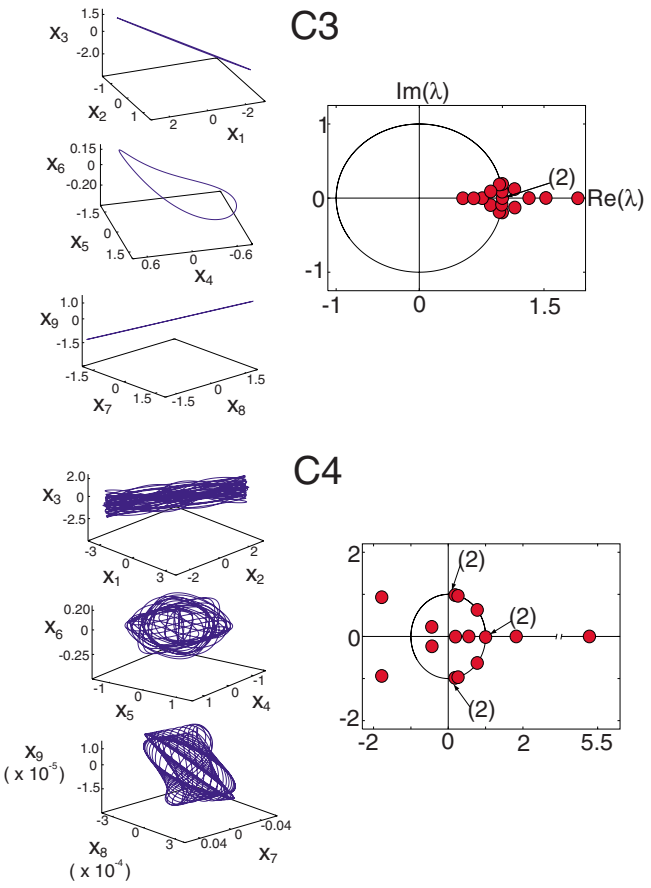

FIG. 10. (Color online) Periodic orbits C3 and C4 of the nine-DOF generalized Hénon-Heiles system.

To conclude, we proposed a full phase-space Newton-Raphson algorithm and a selection algorithm to search for periodic orbits in multidimensional Hamiltonian flows. By abandoning the Poincaré surface of section and incorporating non-Poincaré surface components from the full monodromy matrix into the NR algorithm, we are able to adapt the NR algorithm's iteration process to the natural phase flow of the system. This increases the success probability of our FPSNR algorithm in locating periodic orbits compared to Poincaré surface-based methods because we do not have to deal with additional artificial constraints imposed by the Poincaré surface on the iteration process. Moreover, since it is just a simple Newton-Raphson algorithm, our method is computationally low cost and hence genuinely implementable in high dimensional systems. Also, our method can successfully converge onto a periodic orbit regardless of its stability characteristics. In our numerical examples, we demonstrated that the method is applicable to high dimensional flows with as many as nine degrees of freedom.

ACKNOWLEDGMENTS

One of the authors (Y.W.K.) thanks the Japanese Government (Monbukagakusho: Mext) for support, and Yasuki Arasaki and Tom Gally for discussions. This work has been supported by Grant-in-Aid for Scientific Research from the Ministry of Education and Science in Japan.

APPENDIX A: USING GRIDS IN THE SELECTION ALGORITHM

In the selection algorithm, points are taken by laying a grid in phase space. The spacings between points on the grid

TABLE V. Values of energy, perimeter, and Δ for periodic orbits B1 to B5 and C1 to C4 of the six- and nine-DOF generalized Hénon-Heiles systems.

Systems B and C			
Name of periodic orbit	Energy	P	Δ
B1	14.84	35.73	3×10^{-5}
B2	9.86	206.72	5×10^{-6}
B3	9.88	564.41	2×10^{-6}
B4	15.23	35.06	3×10^{-5}
B5	11.90	30.82	3×10^{-5}
C1	9.78	28.17	4×10^{-6}
C2	11.30	30.20	3×10^{-6}
C3	10.98	29.77	3×10^{-6}
C4	9.89	727.57	1×10^{-7}

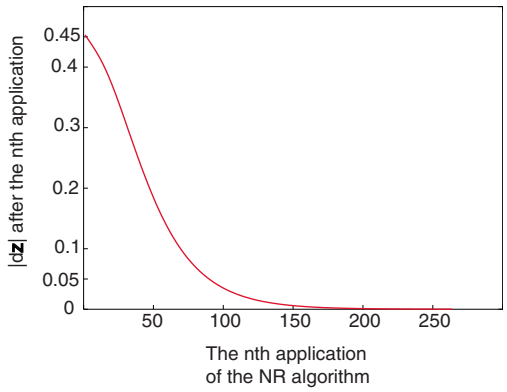

FIG. 11. (Color online) Graph showing the reduction of $|d\mathbf{z}|$ using the FPSNR algorithm of a trajectory leading to the periodic orbit C2. ϵ=0.10 and γ=0.0001. The horizontal axis indicates the nth application of the algorithm, and the vertical axis indicates the size of $|d\mathbf{z}|$ after the nth application. The $|d\mathbf{z}|$ of the trajectory of the initial guess was approximately 0.45, and this was reduced below 0.0001 with fewer than 300 repetitions of the FPSNR algorithm.

can be controlled to manipulate the average volume of phase space occupied by each grid point, thus allowing one to set the size of the ball radius R (Fig. 1). Although employing an extensive grid in the full phase space may seem like a computationally costly exercise especially in high dimensions, one can adapt the grid to specific aspects of the system one is interested in and not necessarily have to check every grid point. The main advantage of using a grid is in the control over the grid point spacings so as to avoid unnecessary repeated sampling of points which are too close together.

When the phase space is low dimensional, such as in the case of systems A and D, we laid out an extensive Cartesian grid. For system B and C, however, because the phase space is large, it is computationally expensive to compute N_{guess} for a grid that covers extensive portions of phase space. Hence we focus our selection on points within a certain energy

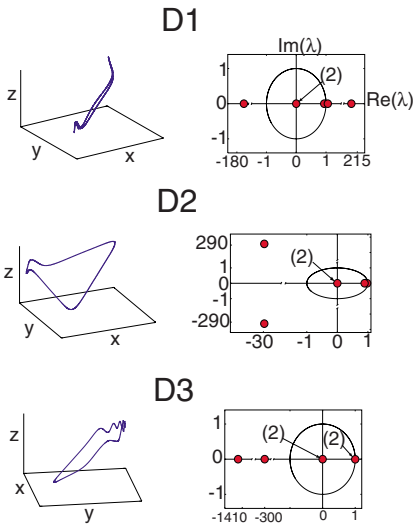

FIG. 12. (Color online) Periodic orbits D1 to D3 of the $x^2y^2z^2$ system. The orbits are three-dimensional curves in configuration space.

TABLE VI. Values of energy, perimeter, and Δ for periodic orbits D1 to D3 of the $x^2y^2z^2$ system.

System D			
Name of periodic orbit	Energy	P	Δ
D1	0.73	22.08	4×10^{-4}
D2	1.77	27.34	4×10^{-4}
D3	6.32	109.15	9×10^{-5}

shell. In our calculations, we simplify matters by fixing the configuration coordinates and making a grid on the surface of a hypersphere in momentum space defined by

$$\sum_{i=1}^{N} \frac{p_i^2}{2} = E_{\text{KE}}, \tag{A1}$$

where N is the DOF of the system and E_{KE} is the kinetic energy of the grid point. To lay out the grid defined by Eq. (A1), we make use of the hyperspherical coordinates. For an N dimensional Euclidean space with coordinates $(p_1,\ldots,p_N)$, the hyperspherical coordinates $(\xi,\psi_1,\ldots,\psi_{N-1})$ are defined as

$$\begin{aligned}
p_1 &= \xi\cos(\psi_1),\\
p_2 &= \xi\sin(\psi_1)\cos(\psi_2),\\
p_3 &= \xi\sin(\psi_1)\sin(\psi_2)\cos(\psi_3),\\
&\vdots\\
p_{N-1} &= \xi\sin(\psi_1)\cdots\sin(\psi_{N-2})\cos(\psi_{N-1}),\\
p_N &= \xi\sin(\psi_1)\cdots\sin(\psi_{N-2})\sin(\psi_{N-1}),
\end{aligned} \tag{A2}$$

where ξ is the radius of the hypersphere:

$$\sum_{i=1}^{N} p_i^2 = \xi^2. \tag{A3}$$

The last angle ψ_{N-1} has a range from 0 to 2π while the rest of the angles have a range from 0 to π. This range covers the whole sphere. The radius of the hypersphere ξ is related to E_{KE} by

$$\xi = \sqrt{2E_{\text{KE}}}. \tag{A4}$$

In our calculations for system B where N=6, we took E_{KE}=15. There are five hyperspherical angles ψ_1 to ψ_5. For ψ_1 to ψ_4, along each angular coordinate 9 points are taken across the range of 0 to π with angular spacing of $\pi/10$ between points. For the last angle ψ_5, 20 points are taken with the same angular spacing across the range of 0 to 2π. This gives $N_{\text{grid}}=9^4\times20=131\,220$. The procedure is similar for system C. E_{KE}=10 and the angular spacing between points along all the angular coordinates is $\pi/6$. This gives $N_{\text{grid}}=5^7\times12=937\,500$. The configuration coordinates for both cases are taken to be at the origin of configuration

space. The grid region, grid spacing, and N_{grid} for the two systems are summarized in Table II.

APPENDIX B: DETERMINATION OF t_{cut}

In the selection algorithm we need to determine t_{cut}, the longest time an integrator can provide accurate calculation for the stability matrix along a trajectory. We used the sixth order symplectic integrator in our calculations [30,31]. When the system exhibits chaotic behavior, machine roundoff errors very quickly accumulate during the integration of the stability matrix and significantly affect its accuracy. As the correction vector $\mathbf{c}$ of Eq. (10) in the FPSNR algorithm is computed using the stability matrix, in order to be sure that the direction of $\mathbf{c}$ before rescaling is reliable, one must ensure the numerical accuracy of $\mathbf{M}'(t)$ when the trajectory makes its closest approach to its starting point.

One common practice for checking the numerical accuracy of $\mathbf{M}'(t)$ is to take its determinant and make sure that it is unity. Here we describe a more stringent criterion based on the symplectic property of the stability matrix. In Hamiltonian systems, the stability matrix $\mathbf{M}'(t)$ satisfies the symplectic condition [33,34]:

$$\mathbf{M}'(t)\mathbf{J}\mathbf{M}'^{T}(t) = \mathbf{J}, \quad t \geq 0, \tag{B1}$$

where

$$\mathbf{J} = \begin{pmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{I} & \mathbf{0} \end{pmatrix}, \tag{B2}$$

and $\mathbf{M}'^{T}(t)$ means the transpose of $\mathbf{M}'(t)$. Any violation of Eq. (B1) signals that significant roundoff errors have been accumulated in the course of the calculation. As a measure of whether Eq. (B1) is satisfied, we define as $\phi(t)$ the largest matrix element of the matrix

$$\text{abs}[\mathbf{M}'(t)\mathbf{J}\mathbf{M}'^{T}(t) - \mathbf{J}], \tag{B3}$$

where abs$[\cdot]$ means to take the absolute value of every element of the matrix in the parentheses. If $\phi(t)$ is zero, then the symplectic condition is satisfied perfectly; deviation of $\phi(t)$ from zero is a measure of the extent of deviation of $\mathbf{M}'(t)$ from the symplectic constraint, and hence of the amount of accumulated roundoff errors.

To determine t_{cut} for each point, integrate the trajectory and its stability matrix over time and monitor $\phi(t)$. When $\phi(t) > \delta$ where δ is a small constant, the symplectic condition is considered as violated and the rest of the trajectory is discarded. The largest t for which $\phi(t) < \delta$ will be the t_{cut} for that point.

APPENDIX C: PARAMETERS

1. δ

The parameter δ in Appendix B is used to set an upper limit on the $\phi(t)$ of a trajectory in order to determine its t_{cut}. For all our numerical calculations, δ is set at 10^{-10}. When the trajectories are strongly unstable (such as in system D or in the high-energy regimes of systems B and C), $\phi(t)$ becomes greater than 10^{-10} very quickly. Consequently, if R is chosen too small in these cases, the selection algorithm may return very little or even no initial guesses at all.

2. R

The parameter R is used in the selection algorithm in Sec. II A to define the radius of a ball around the starting point of a trajectory (Fig. 1). There is no simple formula for determining the appropriate R for use in any arbitrary system. First, without specifying the units the numerical value of R by itself has no intrinsic meaning since only the size of R relative to the volume of phase space under consideration can give an indication of whether the ball is big or small. A good choice of R comes from having some acquaintance with the system. A rough guide to choosing R would be to consider the phase space volume covered by one's grid and estimating the length of one side of the volume; R should then be roughly from $\frac{1}{10}$ to $\frac{1}{20}$ of the length of the side. Alternatively, one can use the spacing between the grid points as a guide, provided one's grid is not too sparse compared to the total volume of phase space one wishes to cover.

For the systems we investigated in this paper, we found that a length of around 3.0 along each dimension usually encompasses a reasonably large volume of phase space, with many trajectories within that volume having energies high enough to dissociate. We learned from experience that a good choice of R for these cases is between 0.1 and 0.5. With R smaller than 0.1, one might miss the unstable periodic orbits; with R greater than 0.5, the selection algorithm may return many choices which are not good candidates for locating periodic orbits, hence consuming computational time unnecessarily.

3. ϵ, γ, and the step limit

The parameters ϵ, γ are used in the FPSNR algorithm in Sec. II B. ϵ is used to scale the length of the correction vector relative to the magnitude of the vector $d\mathbf{z}$ [cf. Eq. (10)]. γ is the tolerance length against which $|d\mathbf{z}_{\text{new}}|$ is compared in order to decide whether to repeat or abort the FPSNR algorithm. We found that a good value for ϵ is 0.05 for most of the systems we investigated. γ should be set at a value such that the error due to $|d\mathbf{z}|$ will not incur significant errors in the quantities (such as the action integral) computed using the periodic orbit.

There is another important parameter not explicitly discussed in the text, which we shall call the *step limit*. It is the number of repetitions the FPSNR algorithm attempts to apply the correction vector before discarding a guess as a bad periodic orbit candidate. If one is applying the FPSNR algorithm to a large number of initial guesses, a practical choice of step limit would be between 200 and 300. If one is focussing on only a few selected initial guesses, then the step limit can be set much higher. In the case of applying the FPSNR algorithm to a large number of initial guesses, one way to cut down on the computational time spent on the FPSNR algorithm will be to reduce the size of $|d\mathbf{z}|$ in several stages. If γ is set straightaway to a small number and one makes the step

limit large, the FPSNR algorithm may spend too much time on some initial guesses which might turn out to be failures. A better strategy would be to first set γ at an intermediate value and use a smaller step limit. After this first stage of reduction, γ is lowered and the FPSNR algorithm is applied again. This way, bad candidates can be eliminated at earlier stages of the convergence process and much computation time can be saved.

[1] J. M. Greene, J. Math. Phys. **9**, 760 (1968).
[2] S. Gekle, J. Main, T. Bartsch, and T. Uzer, Phys. Rev. Lett. **97**, 104101 (2006).
[3] M. C. Gutzwiller, *Chaos in Classical and Quantum Mechanics* (Springer, New York, 1990).
[4] E. J. Heller, Phys. Rev. Lett. **53**, 1515 (1984).
[5] B. Eckhardt, G. Hose, and E. Pollak, Phys. Rev. A **39**, 3776 (1989).
[6] K. Takatsuka, S. Takahashi, Y. W. Koh, and T. Yamashita, J. Chem. Phys. **126**, 021104 (2007).
[7] P. Cvitanović, Phys. Rev. Lett. **61**, 2729 (1988).
[8] C. Grebogi, E. Ott, and J. A. Yorke, Phys. Rev. A **37**, 1711 (1988).
[9] E. Ott, T. Sauer, and J. A. Yorke, Phys. Rev. A **39**, 4212 (1989).
[10] P. So, E. Ott, T. Sauer, B. J. Gluckman, C. Grebogi, and S. J. Schiff, Phys. Rev. E **55**, 5398 (1997).
[11] J. R. Miller and J. A. Yorke, Physica D **135**, 195 (2000).
[12] P. Schmelcher and F. K. Diakonos, Phys. Rev. E **57**, 2739 (1998).
[13] R. L. Davidchack and Y. C. Lai, Phys. Rev. E **60**, 6172 (1999).
[14] R. L. Davidchack, Y. C. Lai, A. Klebanoff, and E. M. Bollt, Phys. Lett. A **287**, 99 (2001).
[15] J. J. Crofts and R. L. Davidchack, SIAM J. Sci. Comput. (USA) **28**, 1275 (2006).
[16] M. W. Hirsch and S. Smale, Commun. Pure Appl. Math. **23**, 281 (1979).
[17] A. Carnegie and I. C. Percival, J. Phys. A **17**, 801 (1984).
[18] V. S. Kalantonis, E. A. Perdios, A. E. Perdiou, and M. N. Vrahatis, Celest. Mech. Dyn. Astron. **80**, 81 (2001).
[19] P. Cvitanović, R. Artuso, R. Mainieri, G. Tanner, G. Vattay, N. Whelan, and A. Wirzba, *Classical and Quantum Chaos Part I: Deterministic Chaos*, Version 11, 2004. Available at http://chaosbook.org/.
[20] S. C. Farantos, Comput. Phys. Commun. **108**, 240 (1998).
[21] Y. Lan and P. Cvitanović, Phys. Rev. E **69**, 016217 (2004).
[22] M. L. Brewer, J. S. Hulme, and D. E. Manolopoulos, J. Chem. Phys. **106**, 4832 (1997).
[23] H. D. Meyer, J. Chem. Phys. **84**, 3147 (1986).
[24] H. Ushiyama, Y. Arasaki, and K. Takatsuka, Chem. Phys. Lett. **346**, 169 (2001).
[25] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in Fortran 77*, 2nd ed. (Cambridge, New York, 1992).
[26] M. L. Brewer, J. Chem. Phys. **111**, 6168 (1999).
[27] M. Nest and H. D. Meyer, J. Chem. Phys. **117**, 10499 (2002).
[28] M. D. Feit, J. A. Fleck, Jr., and A. Steiger, J. Comput. Phys. **47**, 412 (1982).
[29] M. D. Feit and J. A. Fleck, Jr., J. Chem. Phys. **80**, 2578 (1984).
[30] R. I. McLachlan and P. Atela, Nonlinearity **5**, 541 (1992).
[31] K. Takahashi and K. Ikeda, J. Chem. Phys. **99**, 8680 (1993).
[32] S. Takahashi and K. Takatsuka, Phys. Rev. A **70**, 052103 (2004).
[33] R. G. Littlejohn, Phys. Rep. **138**, 193 (1986).
[34] V. I. Arnold, *Mathematical Methods of Classical Mechanics*, 2nd ed. (Springer-Verlag, New York, 1989), p. 219.
[35] See p. 295 of Ref. [19].
[36] If the Hamiltonian involves curvilinear coordinates, $D(t)$ will need to take into account the metric of the phase space.
[37] Of course, there will be a small portion at the beginning of every trajectory where $D(t)<R$; this beginning portion is excluded from consideration.